Fall 2002

LiNK

ARMY RESEARCH LABORATORY • MAJOR SHARED RESOURCE CENTER

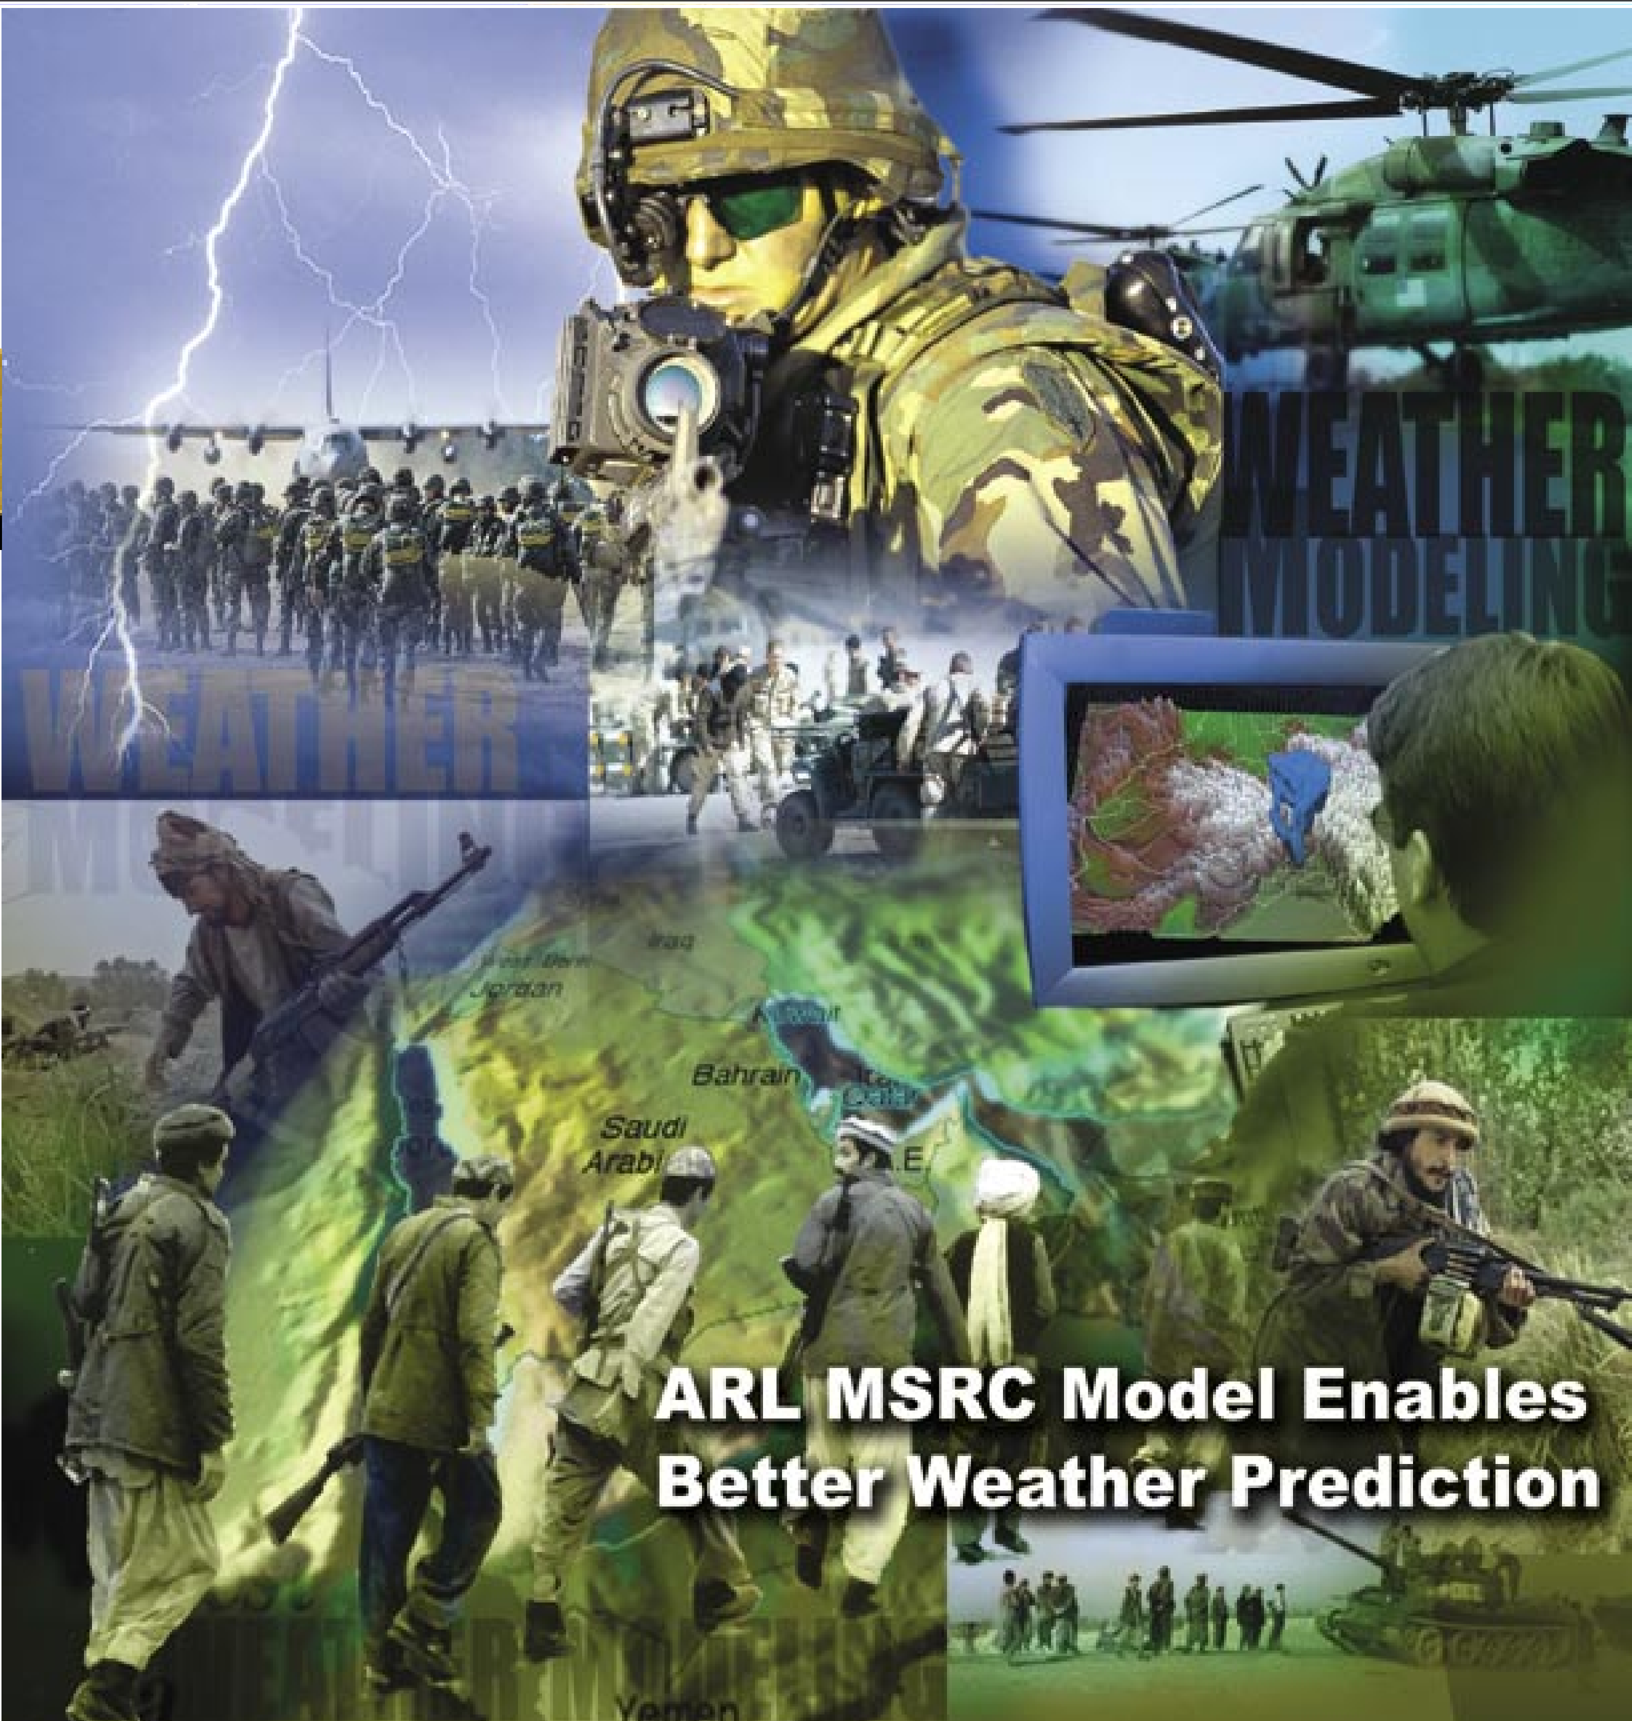

ARL MSRC Model Enables Better Weather Prediction

MSRC Director's Forum
Link Evolves into Journal

Welcome to the new Link. We received feedback during two customer sessions that asked for improved communications. As an initial step toward improved customer service, the ARL MSRC has updated the Web site and transformed the Link. Our goal is to provide you with more information and better customer support by expanding contents to provide application and system support information.

The Link's colorful cover shows recent work in the fight against terrorism. The page count has been increased so we can deliver in-depth technical content, case studies of the latest research projects, attention-getting scientific visualizations, and reviews of the latest technology being deployed at the ARL MSRC. We will continue to interview researchers, describe available resources, profile new work teams, and create an HPC environment responsive to our customers.

As the Center grows to meet customer needs, we will continue to replace aging technology and add cutting-edge supercomputers to our world-class facility. The Hardware Integration and Networking teams set up, tested, and deployed our new 768-Processor IBM Power4-based HPC system in record time. This effort launched the ARL MSRC back into the top ten of the world's most powerful computing sites. Additionally, we were able to offer customers increased computational power for jobs spanning multiple processors.

Recently, the Software Application team merged with Customer Services to give ARL MSRC customers more comprehensive support with integrated solutions. The ARL MSRC has been adding more systems and new customers, which significantly increased support demands. The reorganization will help us make sure we continue to meet and exceed customer requirements. The new Customer Service team lead is Steve Thompson.

To develop methods for holding meetings in virtual space, we are working with PET partners to explore techniques for technical collaboration and conducting online training, consultation, and information sharing. In the future, we will also have Web-based tutorials posted.

The ARL MSRC is also pursuing remote scientific visualization capabilities so remote users can more easily view the results of their computations. In one approach, each animation frame is compressed to reduce the amount of data while maintaining features of interest. The remote users need to transfer just a small percentage of their overall data set to gauge its quality; this is especially useful during preliminary stages of a research project. Another approach uses specialized VideoOverIP hardware so a non-local researcher can view output from an ARL MSRC scientific visualization system in real time on his desktop computer. This approach allows the researcher and a scientific visualization team member to collaborate at a distance to produce high-quality visualization, eliminating the need for travel and time-consuming trial and error.

The ARL MSRC is also expanding communications through exhibits. This November, the Center will host the DoD HPC Modernization Program booth at Supercomputing 2002. The exhibit will describe research areas and program initiatives from the entire HPCMP. We will support this year's SC2002 theme "From Terabytes to Insights" with interactive presentations, lively talks, videos, demonstrations, and posters that illustrate how the DoD uses high performance computing technologies to successfully deliver scientifically based solutions to the warfighter. Special static displays will showcase tangible products developed from research performed with DoD HPC resources. These displays will be located in a special area of the convention center.

All of these exciting changes are part of our focused efforts to have the ARL MSRC provide the best possible HPC environment to our customers. We look forward to working with you as we continue to transform this world-class facility into a communications-rich center. We are committed to fully supporting the DoD computational mission by providing complete solutions to all of your research challenges. We look forward to your comments which can reach us through e-mail: outreach@arl.army.mil.

We will meet again soon.

Charles J. Nietubicz

Mr. Charles J. Nietubicz is the Division Chief of the High Performance Computing Division (HPCD) and Director of the Major Shared Resource Center (MSRC), Computational and Information Sciences Directorate (CISD), Army Research Laboratory (ARL).

In This Issue:
ARL MSRC Enhances Weather Modeling Capability

Signature Modeling of Zsu 23

User's Group Conference 2002 Photo Highlights

Table of Contents

MSRC increases Software Capabilities

During the third quarter of 2002, the MSRC increased from 4 to 6 the Fluent parallel 128-network tokens and added an additional network seat to GASP, bringing the total to 2 GASP network seats. Also this quarter, the latest versions of Abaqus, ICEM, CFD, Patran, Matlab, Ansys, Hypermesh, LS-Dyna, and Lisp were installed.

IBM Power4 Systems Arrive at ARL MSRC

As part of the Technology Insertion 2002 (TI-02) effort, the new IBM Power4 systems had pioneer usage during the last quarter and went into full production on 1 October 2002. One Power4 system has 64 processors and 96 Gigabytes of memory and the other has 768 processors with 768 Gigabytes of memory. Both systems and all 25 nodes of the Power4 have been fully integrated and are now in operation. (See story on page 5.)

Mike Cook unwraps an IBM Power4 cabinet.

Dr. N. Radhakrishnan, CISD Director, speaks to a group of visitors from West Point during a tour of the ARL Major Shared Resource Center.

MSRC Hosts West Point Professors

The ARL MSRC was proud to host a group of 25 professors from the U.S. Military Academy at West Point in August. Dr. N. Radhakrishnan, CISD Director, gave the group an overview of the center and its workings, while other MSRC staff members gave the professors a variety of demonstrations. The group, which consisted mostly of military officers, was given a tour of the computer facilities as well as scientific visualization briefings.

Meyeroff Scholars Get Glimpse of Future

Also this summer, approximately 25 incoming freshmen Meyeroff scholars from the University of Maryland in Baltimore were treated to a tour of the MSRC. These outstanding students, all of whom are majoring in Computer Science, were able to see the latest in supercomputing technology as well as scientific visualization demonstrations and even a virtual tour of the MSRC's computer room.

One of the more interesting demonstrations developed by the U.S. Army was given by SciVis team member Kelly Kirk. This application immerses the user in various simulations of real-life Army situations such as basic training, hostage rescue, and weapons training.

The simulation, which was demonstrated on the Elumens VisionStation, is called "America's Army" and was developed as a promotion and potential training tool for the Army.

We hope that these exceptional students came away from this experience inspired by the important and exciting work the ARL MSRC does. Who knows — just maybe some got a glimpse into their own future here with us! (For more information or to download America's Army, go to www.goarmy.com or www.americasarmy.com.)

ARL MSRC Resource Realignment

With the recent introduction of IBM's newest and most powerful HPC platform, the 768-processor P690 Power4, the integration team at ARL will be reallocating some of ARL's HPC resources to better balance and support the overall userbase. The 512-processor IBM Power3 system will be reassigned in the unclassified environment in late

October. The system, previously known as *BAT*, will be upgraded to a full 1GB per processor and available for use in the unclassified environment under the hostname *Brainerd-n1*. This new system will provide an additional 768 GFLOPs in the unclassified environment and will be available alongside the existing 512-processor Power3 system, *Brainerd*. In a phased integration approach, the ARL Integration team plans to merge these two 512-processor Power3 systems into a single, 1024-processor, 1024-GB system later this year. This new combined system is expected to be available for production use by mid-December and will assume the hostname *Brainerd*. Stay tuned in November for updates on the exact dates for the system merge.

According to the Top 500 Supercomputing Sites Web site, the ARL MSRC is the tenth most powerful supercomputing site in the world. For more information, go to http://www.top500.org/list.

ARL MSRC's Early Experience with IBM Power4 Scalable Power System

By Mike McCraney

The IBM P-Series 690 Power4 scalable parallel system was selected for use at the ARL MSRC as part of the Technology Insertion 2002 (TI-02) effort. The ARL MSRC has installed and integrated two systems, one with 64 processors and one with 768 processors. Each P-Series 690 frame consists of up to 32 Power4 processors and may be partitioned into multiple logical nodes, each with its own processors, memory, disk, and operating system.

The Power4 architecture is revolutionary in that a single chip contains dual 1.3-GHz processors with private level 1 caches and a larger shared level 2 cache. In addition, an off chip level 3 cache is shared among multiple processors.

A study was conducted to provide a precursor look at what serial performance and scalability increases ARL users can expect as they transition from the Power3 system to the Power4. At the time the white paper was written and published, the 64-processor system had been installed and integrated at the ARL. However, the production 768-processor system was being pre-staged at IBM's Poughkeepsie, NY, facility.

CTH was the code selected and several combinations of logical partition sizes and switch adapter counts were compared to the performance in a single 32-processor node in the study. HPCMP is particularly interested in CTH performance, as this code is used in millions of CPU hours each year and a large database of CTH performance already exists, allowing the results of the study to be put into perspective.

The study results indicated that the serial performance of the IBM P-Series 690 is very impressive and delivers slightly greater than a three-fold increase over its predecessor, the IBM RS/6000 SP with the 375-MHz Power3 microprocessor. Scalability within a single node was also seen to be very impressive. While there is room for improvement in the internodal communications performance, the introduction of support for dual switch adapters per node shows significant improvement over the single switch adapter configuration. Since this study, pioneer user results on the production 768-processor system have thus far been consistent with those projected in the study.

For more detailed information, see the complete white paper in the proceedings of the 2002 HPC Users Group Conference held in Austin, TX.

***Source**: "Early Experience with IBM Power4 Scalable Power System" by Thomas M. Kendall and Stephen J. Shcraml, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD; and C. Michael McCraney, Raytheon, Aberdeen, MD.*

In-Office Collaboration and Training Comes to ARL

By Dr. Steven Wilkerson

So one might say, "What's next? I have a really good PC. Heck, my old 450-MHz PC wasn't half bad. I've got plenty of disk space, plenty of speed, great software, and my Internet connection is unbelievably fast. What's more, I can call the help desk for almost anything. Life on a PC has never been better." However, with software constantly changing and systems constantly upgrading, there is still much to learn.

Playing phone tag with people to get questions answered and taking courses can be a real hassle. Sometimes it can even take twice as long to learn new software once it's out. "Is this the result of getting old?" The solution may be right around the corner, and the age of true online collaboration may be at hand! ARL MSRC is teaming with HPTi (the PET prime contractor) to develop software and hardware solutions for a number of these problems.

Collaboration nowadays might consist of phone calls, lots of e-mail messages, trips to "exotic" places like ARL's Adelphi, MD site, and links or pointers to files like */usr/local/confused_user/afterhours/workload/toomuchwork/file_does_nothing.cmd.* Once the user finally finds the system with the link, he may find out that the protection on the file was set incorrectly and the guy he was working with is on vacation for the next two weeks. Then the waiting game begins, and everything gets delayed a little more. In spite of this, Internet connection speed is such that real-time collaboration over the Internet or local intranet is possible using a variety of methods. Unfortunately, there is currently no single all-encompassing solution to collaboration.

Therefore, the solutions are piecemeal. The good news is that when they work, they work really well. For example, NetMeeting is provided free, comes installed on Windows 2000 and XP machines, and is available for other Microsoft (MS) and non-Microsoft operating systems as well. Microsoft NetMeeting allows users to share applications, view one another, and talk in real time. Conferencing Server is the next logical step in the process. Conferencing Server allows clients to work as groups, share applications, and communicate like MS NetMeeting via a Web page.

Multicast-enabled networks allow users to work together in larger groups in real time without leaving their offices.

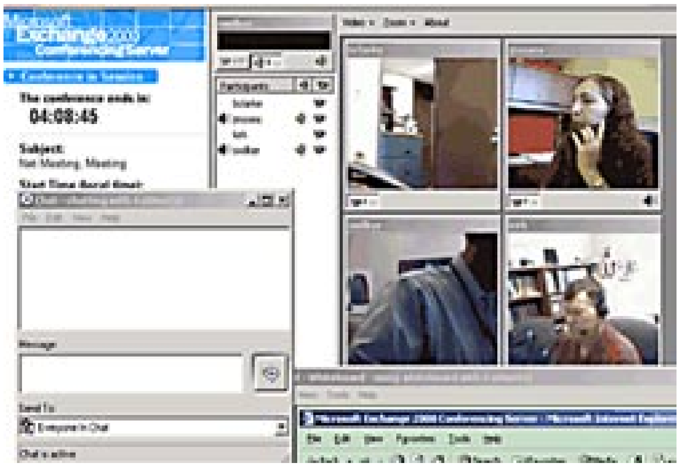

Users can exchange files and messages and share applications in real time rather then looking for elusive links.

With MS Exchange on the way for ARL users, Conferencing Server will be a real asset to managers and workers alike. Security will be addressed with SSL connections, IP blocking, and firewalls. One might ask, "How can we use this for training, technical talks, and distance learning?" MS NetMeeting is a low-quality video feed and is not going to address these issues to anyone's satisfaction. Luckily, there are other solutions available for these applications.

For example, it's not unusual for the scientific visualization team to give a presentation a week to visitors who come to ARL. Many of these visitors only have time to visit the Adelphi site and the MSRC is located at the Aberdeen Proving Ground (APG) site.

Recently, a high-speed connection was established between the collaboratorium at Adelphi and room 200 in building 394 at APG using Amnis NAC encoders and decoders. Using our existing Internet connection, this high-speed, high-quality real-time connection was left on for weeks without any problems. This technology allowed researchers to use our current Internet infrastructure to give their scientific visualization presentations locally from building 394 to ARL visitors at Adelphi.

Recently, guests from the Presidential Scholars Program watched in awe and heard speakers from APG talk on one of the three 8-ft wide screens at in the collaboratorium while the speaker's viewgraphs were shown on another 8-ft screen. The third screen was used for videos and other scientific visualizations. On the APG side, the speaker could view what was on the screens at the collaboratorium while interacting with the audience in real time over an Internet connection. Needless to say, this was a huge success.

Another similar technology has been employed to allow users to view presentations at their offices through their desktop machines. This ARL/CDLT initiative allowed speakers to give their presentations while broadcasting their viewgraphs and images (video and audio) to remote sites via the Internet. The software used was based on Real Systems Player and is available for almost all platforms.

The first presentation test was wrapped in a Web page and served through a proxy server to users around the country. The video and audio signals were broadcast from building 330 to a real server in North Carolina. That signal was linked (embedded) from the proxied Web cast, and users were able to view and interact with the presentation through a text-driven interface for questions and answers on that same Web page. Using this technology, the Web cast can be recorded and saved for future use.

Interested parties need not be present to see motivating talks anymore. Moreover, this same system is being integrated with existing PolyCom sites for presentations and distance learning at ARL and elsewhere. So when asked what's next, the answer might be less travel and more time at home, thanks to Web-enabled technologies and collaborative efforts at ARL!

ARL MSRC Introduces MediaNet Services

By John Vines and Brian Simmonds

In an effort to promote collaboration and cost savings, ARL MSRC has established MediaNet Services. This service encompasses various technogies to allow collaboration over existing networks.

MediaNet services incorporate the use of real-time, compressed audio and video, IP-based communications. IP-based communications are currently being used for video productions, scientific visualization, and remote collaboration.

MediaNet has been providing the ability to collaborate on technical and marketing animations and videos, scientific visualizations, and computer graphics applications without requiring scientists and engineers to travel to geographically dispersed areas across the country.

Objectives:
By leveraging and developing new communications and collaborative technologies, scientists, researchers, and engineers can spend their time and funding on solving problems instead of travel. In addition, the user can interact with the remote person or group at any time rather than make a scheduled trip.

Methodology:
Traditional approaches to collaboration have been to attend conferences or travel to remote research locations to review technical processes and data. Teleconferencing is also being used for multiple researchers/multiple locations to review processes/data, sacrificing video quality for multiple location capability and requiring costly hardware bridging equipment and dedicated connections.

The use of VideoOverIP at the ARL MSRC provides desktop users with DVD-quality video. Using visualization software on MSRC supercomputers, real-time visualizations are generated by a SciVis team member. We then take the computer graphics signal and compress it into an mpeg stream that we transmit via a network connection. During this process, feedback from the remote user to the operator allows the visualization to be optimized and tailored to the needs of the remote user.

Various technologies can be employed to receive the output stream from the encoders. Products range from a software decoder to a dedicated network hardware device.

The ARL MSRC is currently exploring ways that this technology can be implemented. Currently a link between the Aberdeen Proving Ground and Adelphi, MD, visualization laboratories is on full time and is proving to be quite useful for collaboration between the two locations.

For more information, contact John Vines at jvines@arl.army.mil.

What's all this Scripting Stuff, Anyhow?

By Jerry A. Clarke

Back in the "good old days" of computing (circa 1985), an engineer needed to build a computer interface board from combinational logic chips, write a device driver in assembly language, and re-link the operating system before actually getting around to writing an application. Back in those days, my favorite programming language was "solder," assembly language was a necessary annoyance, and FORTRAN was for sissies.

But today I am no longer a knuckle-dragging Neanderthal who scoffs at higher-level programming languages. No, I am now a sophisticated connoisseur of computing, who develops object-oriented interfaces for scalable computing platforms that I've never actually touched, in languages that didn't exist in 1985. This amazing metamorphosis can be traced to a single line of code :

```
button .b –text "Hello" –command "puts Hello"
```

This single line of Tcl/Tk code defines a pushbutton with text that will print "Hello" when activated. It works on MS Windows, Macs, and virtually every UNIX platform known to man. Having just finished developing a user interface in C under X/Motif when I saw this, I immediately realized the potential for time savings and portability. Digging a bit deeper, I found other advantages to these higher level "scripting" languages that are less obvious but perhaps even more dramatic.

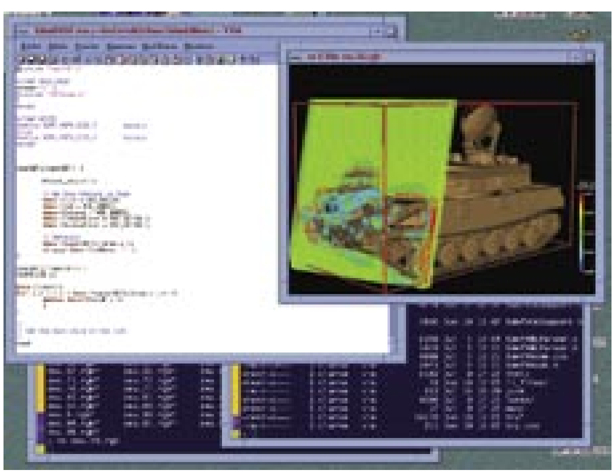

So what is a scripting language and why should you use one? Well, to clearly answer these questions, let's first look at some other types of languages. Above "solder" and machine instructions, there is assembly language. Each line of assembly, roughly equates to one machine instruction. Handling argument passing to functions is an exercise left to the user. Since each type of processor has its own assembly language, portability is an "issue" to say the least. The C programming language somewhat alleviates the portability problem. C is basically a machine-independent assembly language that hides the details of many common tasks like subroutine access and conditional looping (while, for, etc.).

A single line of C may result in multiple machine instructions, but it's typically less than 10, so it's still very low-level programming. C allows for the use of "pointers," which gives programmers the power to shoot themselves in the foot. So many scientific programmers, who wish to express an algorithm without worrying about the underlying machine details (sissies!), choose more structured system programming languages like Fortran77. Fortran77 compilers have an easy time optimizing the machine instructions well beyond the capability of all but the most determined assembly language programmer. Other languages add quirky syntax to introduce more programming abstractions like "object-oriented" programming (C++) and array operators (Fortran90).

These languages, however, have two main characteristics in common. First, they are compiled. That means that the human readable source code is transformed into executable machine instructions. Second, they are what John Ousterhout, the inventor of Tcl, calls "strongly typed." In other words, variable type is defined during compilation. Now, neither of these qualities is bad. In fact compiling strongly typed code catches many mistakes before runtime, for example passing a float instead of an integer. These languages are perfect for implementing computationally intensive algorithms from scratch.

Scripting languages like Perl, Tcl, and Python, on the other hand, are weakly typed. For example, the Python command print MyVar will print the value of MyVar whether it's an integer, float, or even a character string; the "type" is determined at runtime. In addition, these languages are either interpretive or compiled to an abstract "byte code" as opposed to real machine instructions. This allows programs to easily construct commands during runtime. For example, the Tcl/Tk code for a user interface can be constructed by the program itself, and then executed. These features, however, result in code that typically is less efficient. That's not an issue in a user interface but is very important in a loop that will be executed millions of times.

To bridge the gap, scripting languages like Perl, Tcl, and Python allow for the "wrapping" of system programming languages. In other words, you bundle your favorite, efficient subroutines, along with some interface code, into a shared library. This library can then be loaded at runtime, and accessed from the scripting language. A great example of just how powerful this can be is the Visualization Toolkit (vtk). It's a huge C++ class library of scientific visualization and graphics routines that has been wrapped for Tcl, Python, and Java. There are several freely available pieces of scripting code that will query the object's methods at runtime and produce a graphical user interface (GUI). In other words, they can produce a GUI for an object that they have never seen before! Try that in Fortran.

So, system programming languages are great for producing efficient implementations of time-critical functionality from scratch. Scripting languages are great for gluing together large chunks of functionality. Mixing the two is like having a box full of tools verses a swiss army knife. The trick, however, is to choose the right tool for the job. If you're implementing a finite difference solver, you'll probably use Fortran. If you want that solver to gracefully handle free form input from a TCP/IP socket, you might call that Fortran from a layer of C. If you want a GUI for the whole thing, you might "wrap" it for Tcl.

On most of our recent development projects, we've found restricting ourselves to one programming language impossible. We are also finding a great benefit in "gluing" together big chunks of functionality with scripting languages. Solvers written in Fortran can be combined with visualization written in C++ and I/O routines (like HDF5) written in C. We simply provide a scripting interface for each of these chunks in a centrally accessible location and then write a short script to provide a customized tool.

For example, the Finite Element Model (pictured opposite bottom) is stored in HDF5 and described in XML. A Python script parses the XML, reads the HDF5, visualizes the result, and allows the user to rotate, translate, and scale – all in less than 50 lines of code. This is possible because each chunk of functionality has been separately "wrapped." We combine the functionality of several different tools, written in different languages, to produce exactly the "application" we need. And because the scripting languages are portable across architectures, we only need to provide each of the separate chunks of functionality on each platform and all of the "applications" come along for free!

Scripting languages are also finding their way into scientific computing itself. Scalable, parallel solvers are implemented in Fortran or C++ and wrapped into a shared library. Python is then launched as an MPI application and loads the necessary functionality at runtime. SpaSM, a molecular dynamics code, and VTF, a coupled fluid-structures code, are just two examples. This software architecture makes it easy to integrate user written modules and add custom functionality to a simulation.

So the next time your soldering iron won't heat properly, or if you're just tired of being teased by your colleagues for programming in Fortran, give scripting languages a look. They're portable, they're free, and they'll make you a better person.

For a great overview of scripting in general, see the article by John Ousterhout (the orignal author of Tcl/Tk) "Scripting: Higher Level Programming for the 21st Century" at

http://www.tcl.tk/doc/scripting.html

Cover Story:

Scientists Exploit ARL's High Performance Computers to Enhance Weather Prediction Tools

By Pat McKenna

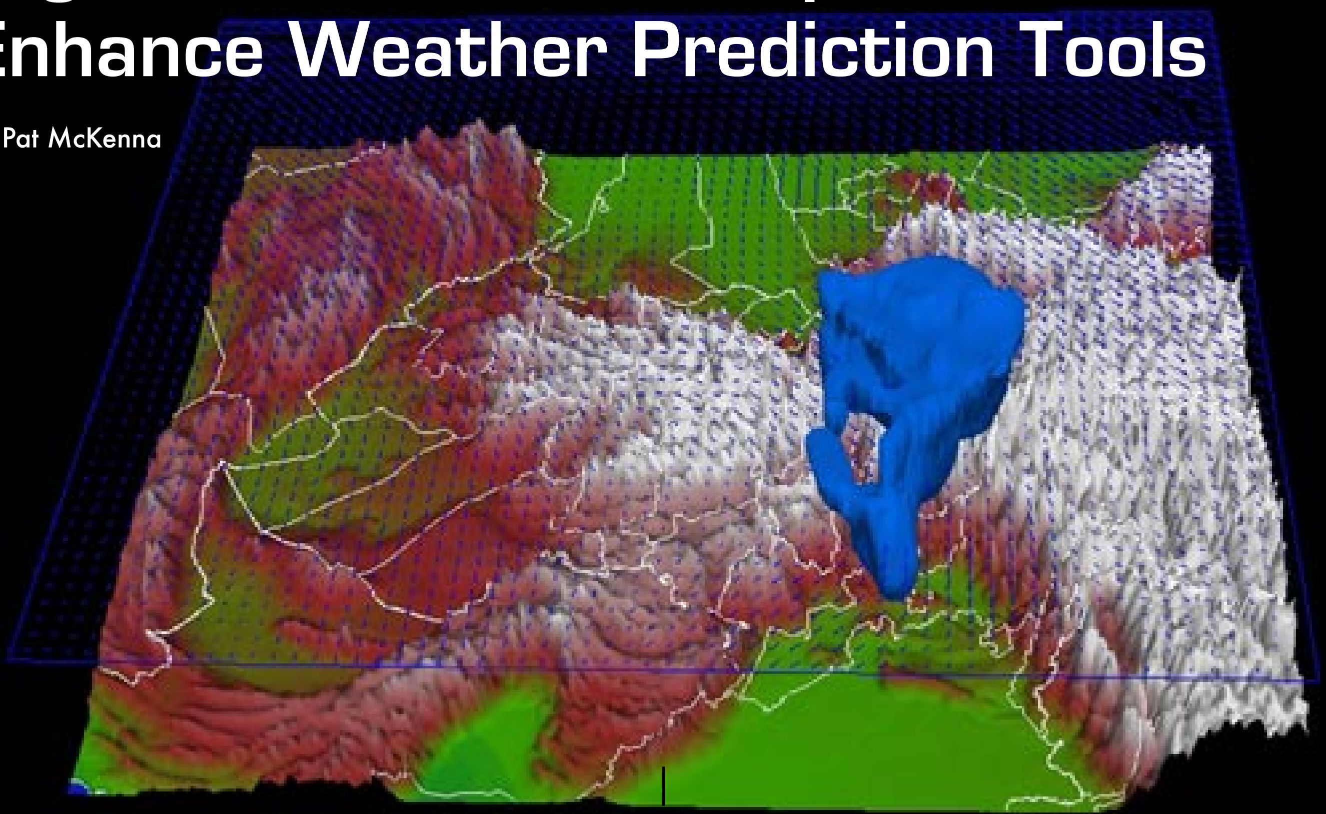

9km resolution run over Afghanistan. The wind vectors are plotted approximately 7.1km above sea level with the clouds in blue.

Mathematicians don't back down from dares. Just double-dare one to name the first 50 digits of π and then stand back! So when the Army Research Laboratory's Computational Information Sciences Directorate (CISD) was given the challenge to forecast weather at horizontal resolutions of one kilometer, it was like betting on the probable rise of the sun in the east.

Chat Williamson, a mathematician in the Directorate's Battlefield Environment Division, rose to the challenge to see if the ARL Major Shared Resource Center's (MSRC's) supercomputers could help better predict the ever-changing weather patterns that can affect military operations. With better weather forecasting, military mission planners can take some of the guesswork out of battlefield planning.

Military operations are often performed in lands of extremes. These areas often boast high mountains and sharp peaks with unforgiving, bone-chilling winter weather and deep freezes. In contrast, these same areas may be bordered by flat, barren land with desert-like conditions consisting of sizzling summers and blinding sandstorms. This creates quite a challenge for military troops performing maneuvers in this varying environment. If military weather forecasters were able to better predict the local weather in these areas, then planners could keep soldiers, airmen, sailors, and Marines safer by avoiding routes that take troops through bad patches of weather.

Unfortunately, the models that most weather forecasters use are at a 10-kilometer resolution, which, given rugged terrain and varying climates, is woefully deficient.

"Everything that is happening below 10 kilometers is estimated. It's an average," said Chatt Williamson, a five-year veteran of ARL and a frequent user of the ARL MSRC. "A lot of different meteorological phenomena occur under 10 kilometers that just get missed." Translation: Under 10 kilometers, weathermen make an educated guess based on more general information.

Because of the complexity of the problem, Williamson exploited the massive computational power of the ARL MSRC to see if he could boost the resolution of the weather models from 10 kilometers to one kilometer over the varying terrain of Afghanistan.

Williamson worked with Dr. Greg Tripoli, a meteorology professor at the University of Wisconsin, using the University of Wisconsin Numerical Modeling System (UW-NMS), also known as N-BFM, the Nonhydrostatic Battlefield Forecast Model. Meteorologists in the Washington D.C. area and in Madison, WI, use this model to make their daily forecasts.

After a week's worth of computer runs, using 192 processors of the MSRC's SGI Origin 3000 supercomputer, the scientists were able to achieve a horizontal resolution for better than real-time computation at three kilometers. The model's domain covered almost all of Afghanistan, except for a small part of the Hindu Kush. The model's grid measured in at 500x384x35, and the team was able to achieve four to one under real time. Translation: The supercomputer took six hours to generate a 24-hour forecast at 3K resolution.

Although Williamson was unable to reach the 1K goal, he was still pleased with the results.

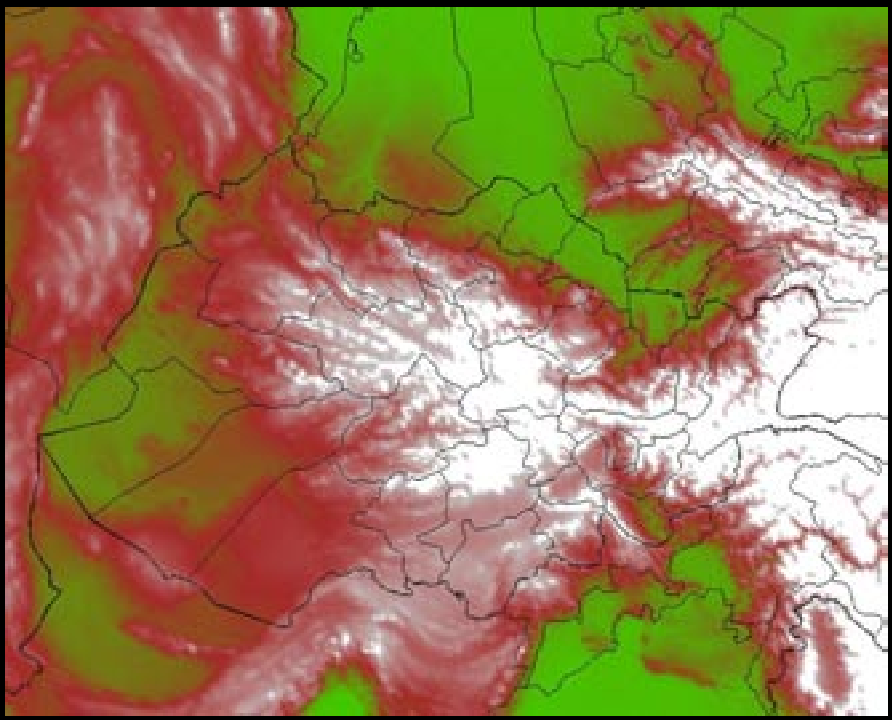

Overhead view of the model domain. Note: Afghanistan has been rotated 35° clockwise.

"To reach the same achievement for the one-kilometer challenge, I surmise that it would take about 27 times the computing power." Williamson said.

Williamson initialized his model from the Global Aviation Model, which the National Weather Service issues. The ARL team used both the NWS's data and model. The team modeled temperature, humidity, cloud cover, surface precipitation, winds, and barometric pressure, among other variables.

"We also looked at turbulent kinetic energy, or turbulence," said Williamson, "which is good to know if you're operating rotor aircraft."

Williamson said that the MSRC's supercomputers were invaluable to his research. In addition, he said that, in essence, he was computing about 100 times the information in the same amount of time.

According to Williamson, if you ran the same model, which took six hours at the MSRC, on a four-processor desktop SGI Onyx machine, it would take about 2000 hours to compute the 24-hour forecast. And on a single-processor Windows-based PC? You might as well read *War and Peace*.

Warfighters need more than just a good guess if they're going to survive in hostile conditions. Williamson's model gives them that fighting chance.

"If the weatherman says there's a 50 percent chance it will rain, we could – using the MSRC's supercomputer and our model – now say there's an 80 percent chance of rain because of our ability to resolve the weather model at a higher resolution," Williamson said. "At 3K resolution, you could better predict the chance of precipitation. The more information you have the better.

"We may not have lost UAVs [Unmanned Aerial Vehicles] or other assets if we had been better able to predict the weather," the mathematician said. "On one of the runs we conducted, we explicitly predicted water content, temperature, and freezing levels. We were then able to visualize these little pockets of super-cooled water and show where you wouldn't want to fly a UAV or helicopter."

"This kind of data could save lives," Williamson said.

Project Team: Chatt Williamson, ARL; Robert Dumais, ARL; Dr. Greg Tripoli, UW-Madison; Tom Kendall, ARL MSRC.

2002 User's Group Conference a Hit in Austin

The User's Group Conference this year was held on June 10^{th} through 14^{th} in Austin, Texas. The ERDC MSRC folks, as sponsors, put on a really good show. Despite the 90+ degree weather at 8AM, the conference was successful. ARL's participation this year was strong with 16 papers presented. Two researchers also participated in the poster session. The social event for the conference was a trip to the Texas History Museum. This evening event featured an IMAX movie, tours through the museum, a multimedia presentation about the history of Texas, and an authentic Tex-Mex dinner.

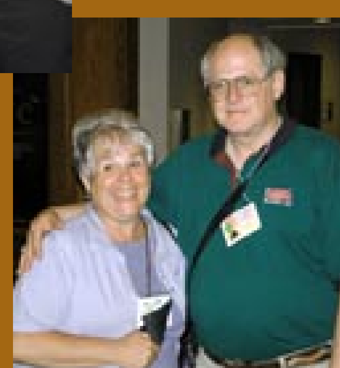

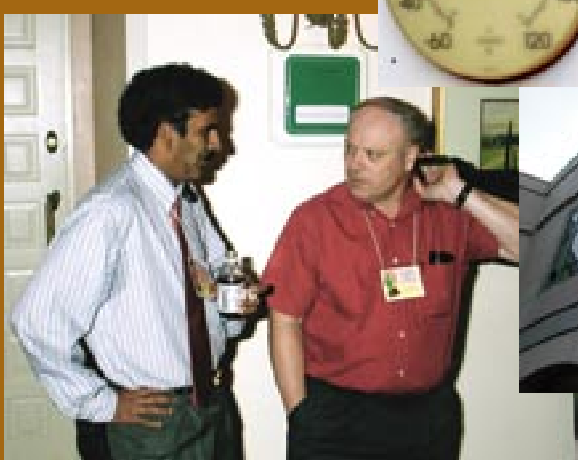

Photos by Brian Simmonds and ASC Outreach Team

Summer Interns at ARL: A Proud PET Tradition Continues

By Robin M. Liles-Meyer

In 1998, the Programming Environment and Training (PET) program developed a residential summer-long internship program for undergraduate and graduate students in mathematics, engineering, and science to encourage young Americans to consider computer science and engineering careers within the Department of Defense and elsewhere. In this program, the students are paired with scientists and engineers at each of the four MSRCs to conduct real-life research projects in the Computational Sciences. The program also strives to include significant participation from Minority Serving Institutions (MSIs).

Right to left, Timothy White receives instruction from his mentors Dr. Charles Macon and Rick Angelini.

This year eight students, one from as far away as Puerto Rico, participated in the eight-week Summer Intern 2002 program at the ARL MSRC. Students were selected by a panel of ARL reviewers based on grade point average, level and type of courses taken (i.e., more credit was provided for challenging courses), and recommendation letters from faculty. The program, coordinated by High Performance Technologies, Inc. for the PET program, began on 10 June with one week of classes and ended 2 August with each student giving a formal PowerPoint presentation of his/her work. In the interim, the interns and their mentors worked together to formulate and execute an Individualized Development Plan (IDP), ensuring they are able to attain new skills and providing an overall valuable experience for the students. Projects this year included such challenging topics as HPC benchmarking, cluster architectures, and Monte Carlo simulations.

Interns also attended seminars and classes to familliarize them with the MSRC and HPC operationns.

Winzer Carty, a senior from Radford University majoring in Information Systems, was guided by his mentor, Dr. Kofi Apenyo, to design, implement, and populate an intelligence-gathering data warehouse that operates in a distributed database environment. Winzer concluded that the data warehouse can be used by upper-level management to analyze data – in particular, data related to terrorism prevention – after the information has been entered from any point on the network. Winzer was enthusiastic about his summer experience and commented, “I was able to design, implement, and populate a data warehouse using a distributed database system, but most importantly it gave me an experience with technology that few people have access to.”

Brandon Hill, a graduating senior from Old Dominion University majoring in Computer Science, worked on the Smart Ordnance Autonomous Reconnaissance (SOAR) project with his mentor, Dr. Steve Wilkerson. Brandon was able to incorporate realistic background scenes and topography into the flyby animation, which significantly enhanced the graphical representation for the simulation package. He completed a set of animations that are ready to be integrated into a full video. Brandon feels that PET’s greatest assets are the experts working for PET. He said, “I was able to meet and question field experts that I couldn’t meet elsewhere.”

Nwokedi Idika, a sophomore at the University of Maryland at Baltimore majoring in Computer Science, worked with both of PET CCM on-sites, Dr. Charles Cornwell and Dr. Mark Zottola. Nwokedi used a Monte Carlo routine coupled with the Jarvis-Patrick algorithm to identify and evaluate the clustering of various conformations. The centroids attained from these clusters are instrumental in the search for the elusive unique molecular structures that elicit the best biological responses. He very much appreciated the extra time and guidance from his mentors for his challenging project. Nwokedi also followed up with program coordinators to suggest additional involvement from his university and its Meyerhoff Scholar

Program for minority graduate students.

Patrick Mallory, a junior from Alabama A&M University majoring in Computer Engineering, was assigned to David Grove. Patrick and Mr. Grove worked on "Benchmarking and Validating an ARES Code." ARES is a computational structural mechanics code still under development at Cal Tech. He implemented ARES for the simulation of the Taylor impact problem, and his results compared favorably with known exact solutions. Furthermore, he carried out benchmarks to obtain scaling factors for ARES using the Taylor problem on the MSRC Origin 3000 and the IBM SP3 systems. During the process, Patrick had to become proficient in the use of numerous pre- and post-processing software programs, such as CVS, Femap, and Tecplot. Patrick remarked effusively that Mr. Grove is a great teacher.


Tonima Mohiuddinis and her mentor George Petit discuss custer computing.

Joel Todman, a graduate student at Morgan State University majoring in Electrical Computer Engineering, was teamed with Jeanne Angelini for his mentor. He focused on network operations, in particular setting up the management of the Catalyst 4006 Switch, which is a network switch located at ARL. He believed that the intern experience has provided valuable material for his thesis on the performance of TCP/IP over wireless links and felt that the best part was "working in a team environment."

Cesar Vidal came to us all the way from Puerto Rico, where he is a sophomore at the University of Puerto Rico Mayaguez Campus, majoring in Electrical Computer Engineering. Cesar and his mentor, Mark Thomas, worked on a project titled "Network for Effective Data Transfer." Under the guidance of Mr. Thomas, Cesar developed a client-server model that will operate on a TCP protocol in order to correct the loss of information that occurs between the client and server while using the current UDP. He then implemented simulation programs using the client-server model, and found that the simulations ran more smoothly, without error in the frames or data collection from the server by the user/client. The ongoing activities at the MSRC impressed Cesar and he observed, "I have seen how the ARL works: with the large amount of diverse projects going on, they still keep on track with every one of them, and integrate and combine those that have things in common."

Timothy White, a graduate student from Jackson State University majoring in Computer Science, was indeed fortunate to work with two mentors. Dr. Charles Macon, PET CEA on-site, was Timothy's primary mentor. Together they worked on a project entitled "EMCC Web-Based Conference Registration System." Timothy worked to improve the code documentation, corrected errors in code, modified routines that produce the output files, and made the Web site more user friendly. He also enhanced the documentation of the HTML and Perl scripts by adding detailed comments. Timothy also worked with Rick Angelini, Scientific Visualization Team Leader for the ARL MSRC. Mr. Angelini guided Timothy on the creation of a new flipbook animation of the electronic component of a missile. This animation shows the time dependent stresses on an electronic component during a duty cycle. Like the other seven interns in the program, Timothy was delighted with his 2002 PET summer internship and commented, "This internship has increased my computer skills, and this allows me to be more marketable."

Tonima Mohiuddinis is a senior at City College of New York majoring in Computer Science. Tonima and her mentor, George Petit, worked on a project that investigated the architecture of the IBM Netfinity Intel cluster (Franks) and the Compaq Alpha Cluster (Powell), and evaluated their performance through benchmarking. For matrix-matrix multiplies, she concluded that the Compaq Cluster had better linear scaling up to 16 processors than the IBM cluster. When asked about her experience with PET, she replied, "I really worked at ARL, one of the largest research facilities in the world; I interacted with real engineers, real data, real problems, and real results. If asked whether I would do this all again, I would surely reply 'yes,' and I would highly recommend the program to others."

Although many interns started the program with some feelings of intimidation about coming to a world-class laboratory ("how can I fit in?"), these feelings dissipated quickly with the confidence gained from hands-on experience and guidance from their mentors. It is clear from their presentations that the PET summer interns have fulfilled their project responsibilities and provided valuable contributions to the program and ARL. Furthermore, they have left lasting impressions with their mentors, who played a significant role in their professional development. In the end, all eight interns considered the program a success and inquired about future opportunities in the Department of Defense.

Computational Electromagnetics

Scalable Pre- and Post-Processing Approaches for Large-Scale Computational Electromagnetic Simulations

By Pat McKenna

Computational electromagnetic (CEM) simulations of full range military vehicles play a critical role in enhancing the design and performance of combat systems including the Army's Future Combat Systems. Modeling of full range military systems involves generating large-scale meshes, solving equations, and visualization or analysis of results in the range of billions of unknowns or grid points. Hence, the overall objective of this project is to develop and demonstrate a scalable CEM software environment to address accurate predictions of radar cross sections (RCS) of full range armored vehicles with realistic materials and complex geometric configurations. A software environment consisting of scalable pre-processing, post-processing, and accurate CEM software is needed to achieve signifigant reduction in overall simulation time for large-scale military applications. In addition to RCS, this high fidelity software environment can also address wide band communications applications.

Recently, researchers from the Army Research Laboratory, developed a scalable computational environment for large-scale CEA applications to exploit the computational power of the ARL MSRC consisting of (a) scalable grid generation based on implicit surfaces and voxel methods, (b) scalable finite difference time domain method, and (c) parallel visualization using network distributed global memory approach.

Dr. Raju Namburu and Jerry Clarke from ARL and Eric Mark, a Senior Systems Engineer from Raytheon, undertook the project to study parallel pre- and post-processing approaches for large-scale computational electromagnetic simulations. For their experiment, they simulated an 8.2- to 9.2-GHz planar wave passing over a Shilka ZSU-23-4 anti-aircraft vehicle, a Russian armored vehicle used to shoot down planes.

The ARL researchers began with a surface model, what Clarke called "an unstructured blob of triangles," of the armored vehicle. They then constructed a structured grid, a box, which encompassed the entire vehicle. They divided the box into smaller cells, running a code that initiated the electromagnetic pulse wave and calculated the effects of the wave as it traveled.

"That gave us a tremendous amount of data, which we output into a visualization showing the effects," Mark said.

Sounds simple enough, right? Yet, the ARL researchers put the MSRC's supercomputers through their paces. The researchers made three to four computers runs with each one taking about 33 hours.

Their computations generated more than 2.4 billion cells, consuming more than 2 Terabytes of volumetric output. A Terabyte, or 10^{12} bytes, is an enormous amount of information. For instance, the entire Library of Congress is about 10 Terabytes.

The team conducted their pre-processing on the MSRC's 512-processor SGI Origin 3000 and then further processed the data on the 512-processor IBM Power3.

"We would've never been able to run this calculation without machines the size of the MSRC's IBM," said Clarke. "You don't come by 512-processor IBMs every day. If we had do this input mesh by hand, it would've taken potentially months to complete.

"We did it at the MSRC in hours," he said.

By identifying a vehicle's hotspots, areas that reflect the most radar energy, designers can reduce the vehicle's profile on radar screens.

"It will let designers do a lot of 'what ifs' and change their design on the fly. It could also help designers see how changes in their design are going to affect the radar signature and the vulnerability of the weapons system," Clarke said. "A small change might have a big effect. We now have a semi-automated tool for doing these calculations. With the perspectives that we can now glean from these models, we can see exactly what features of the vehicle are causing problems in the design or aiding in the design. It can lead to further design changes later on down the road, telling you what things you want to put into the vehicle and what things that you want to stay away from."

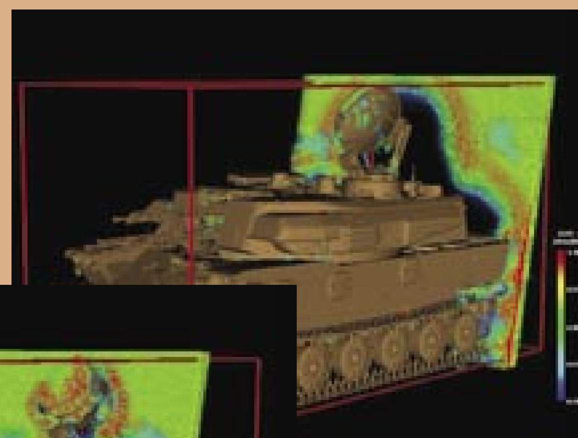

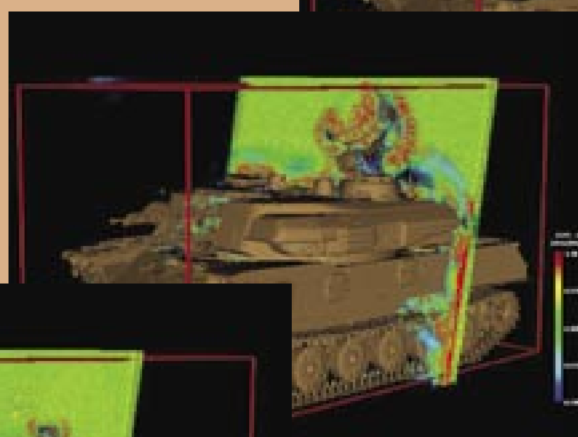

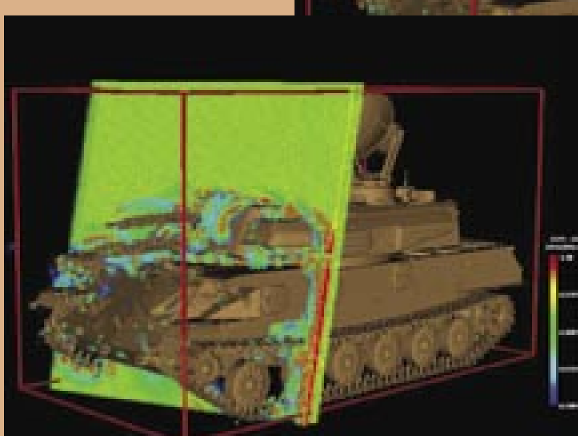

Propagation of an electromagnetic pulse over a ZSU self-propelled anti-aircraft gun. The color represents the magnitude of the total electric field.

What makes the ARL team's model even more useful is that it can be used in an Interdisciplinary Computing Environment (ICE), combining different scientific disciplines during the same simulation. This is not only a more efficient use of resources, but the synergy of the different disciplines gives weapons systems designers the big picture.

"It's like a plug-and-play environment," Clarke said. "Researchers can come along with unique problems and then pick and choose codes they know work, instead of reinventing the wheel and writing the code from scratch. This will greatly increase turnaround time in solving problems, and ultimately, it will allow us to bring better weapons to the warfighter."

With the MSRC's new IBM Power4 supercomputer, the team believes they will be able to solve even more complex problems in less time with the new hardware. The team plans on improving the model's algorithms, testing it with a wider variety of geometries, and streamlining the visualization process. The team believes that their processing approach using parallel supercomputers to create electromagnetic simulations shows great promise.

Who knows what type of stealth technology might be possible with the aid of the MSRC's next generation of supercomputers…

Shock Physics is Subject of Schraml's Research

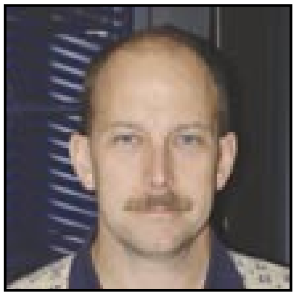

Name: Steve Schraml

Title: Mechanical Engineer

Organization: Weapons and Materials Research Directorate, U.S. Army Research Laboratory

Area of Expertise: Shock physics modeling and simulation. Use of large-scale continuum mechanics codes for modeling problems in transient phenomena such as blast, impact, and penetration.

Project summary:
Development of technologies for enhancement of survivability and lethality of Army ground vehicles; advanced weapons system ammunition, missiles, warheads, armor technologies, active protection systems, novel materials, lightweight materials, and so on. Our goal is to impact the technology development for the Future Combat System.

Impact on DoD:
Did you use any new processes or do something that has never before been done?
We are using finite element and finite volume methods for shock physics that have been in use for a long time. What is new is the ability to run these on scalable architectures. We can dramatically increase the size and complexity of the problems that we can solve. The grids are more refined and that allows you to pickup structural detail that you could not before. In the past when we had significantly less computing power, you would look at a configuration and you would make gross approximations, and only model the most important parts of a weapon system as opposed to an entire system. If you were modeling a missile, you would only model the warhead or the penetrator inside; now we are modeling the whole missile. Then we can consider the whole system instead of the main components in terms of lethality.

Will your project help the DoD cut costs or speed up a process or processes?
Cost savings are realized by the optimization of a system or concept through a combination of numerical simulation, experimentation, and analytical modeling.

How will your project affect end users/warfighters?
We are developing all of the ammunition that goes into the large caliber cannons that are carried on Army ground vehicles, so we do have a direct impact on the warfighter. On the survivability side, the armor development that is done here eventually gets into the vehicles in the field.

Objectives:
What made you decide to start your project?
We do not have a starting or stopping point. Large-scale computational modeling is an integral part of the culture of what we do. We do experimentation, analytical models, and computational modeling as part of an ongoing effort. Using the HPC resources here [at ARL] helps us to accomplish those goals.

Methodology:
What research did you perform?
We look at novel configurations, doing parametric studies to narrow down the number of possible configurations to experiment with. We can eliminate the ones that might not show as much promise and focus on the ones that do. These are the ones that we do the actual experiments on. We don't do fewer experiments, but we do better experiments.

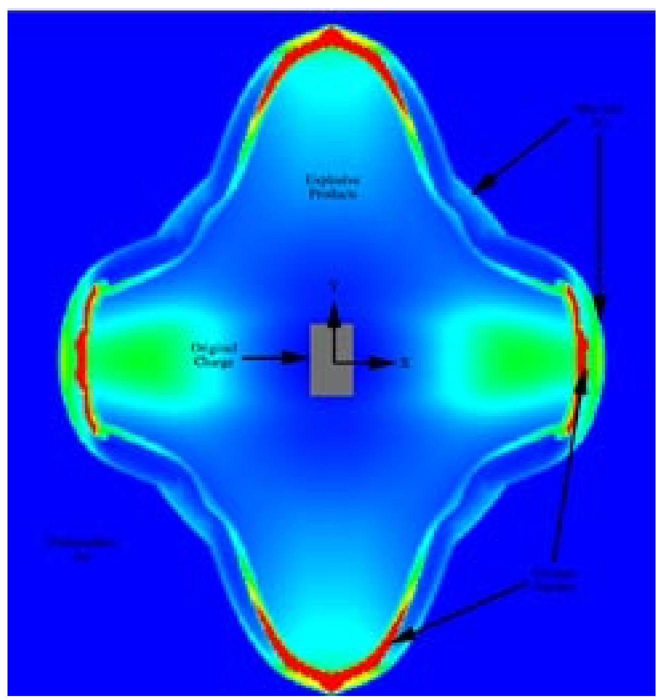

Blast field from a center-detonated cylindrical explosive charge.

Results:
What happened?
On a regular basis we achieve breakthroughs in making ammunition more lethal and lighter weight, reducing the logistical burden so that the vehicles can carry more.

Even if you do not increase the lethality of the ammunition, if you can make it lighter and more compact while maintaining lethality, you allow the soldier to carry more making him more lethal. On the survivability perspective, as we continue to improve the armor systems and make them lighter, it reduces the logistic burden of the vehicle. A lighter vehicle burns less fuel allowing greater range and reduced logistic supply trail.

Significance:

What is the impact of your project?

Readers should care because we have a direct and visible effect on the warfighter. All of the U.S. Army ground combat systems that are in use today are somehow affected by the work we do here and have been for the last 50 years or more.

ARL MSRC Resources: Hardware

What MSRC hardware played a key role in your project(s) and how much did you use the systems?

We primarily use the IBM Power3 512-processor system. Typically we use at least half of that system for a single simulation, sometimes more. We have on occasion had problems big enough that they required the entire 512-processor system. We also use the SGI Origin 3000 256-processor system. We typically use 64 to 128 processors on that system. A typical simulation can consume 10,000 to 50,000 CPU hours. Our group's total HPC usage for FY02 was 5 million CPU hours. We make use of the SciVis assets and the SciVis team to post process and visualize the large amount of data we generate to help make informed decisions.

Software

What software was used for your project?

The primary code we use for simulations is called CTH. It is a finite volume shock physics code developed by Sandia National Laboratories. We worked with Sandia to parallelize the CTH code through a CHSSI project. Sandia did the code development and ARL did the verification/validation and rated the performance. We also use a code called PRONTO3D, which is a finite element code for structural mechanics. Another is ALEGRA, an Arbitrary Lagrangian/Eulerian code. Each code has its own strengths and weaknesses. We try to have the most diverse toolbox possible to allow us to apply the best tool to a given problem. Sometimes a problem doesn't fit into the strength of any one of the codes, so we run the same problem through two codes (finite element and finite volume) to make sure we are getting realistic results. To do our visualizations, we use Ensight along with DICE/ICE.

Did ARL MSRC supercomputers make a difference in your project?

Yes, without HPC resources our project just couldn't be done.

Personnel:

Which ARL MSRC team members helped you with your project?

We receive a lot of support from Tom Kendall, Kathy Smith, and George Kraft, as well as many of the other administrators. We work closely with the system administrators to configure the operational parameters (tuning) of the system to get maximum throughput for our applications. We also get a lot of help from the SciVis team, Rick Angelini, Eric Mark, Jerry Clarke, who help us do our visualizations.

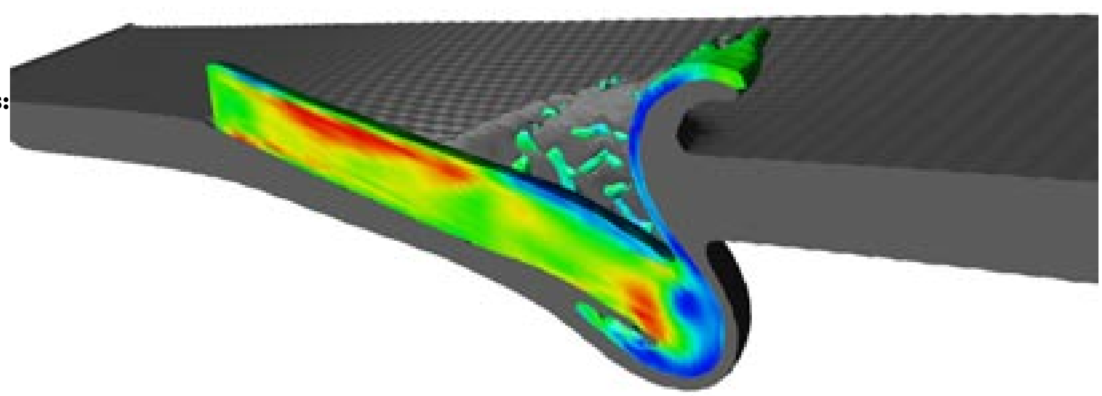

Oblique, long-rod penetrator interaction with a target plate.

Other projects:

What other projects have you worked on using ARL MSRC resources?

The HPC Program looks at our group as one project, but that doesn't mean that we are working on one thing. We are working on a variety of technologies all under the blanket of shock physics, impact, penetration, detonation, explosives, and others. We have a variety of goals under this single HPC project.

What career achievements make you feel proud?

We were one of the winners of the 2001 Army R&D Achievement Award for work that we did on lightweight kinetic energy missiles. That was a combined project where we used experimentation, computational modeling, and analytical modeling.

Editor's Note: *In an effort to correlate a human element to the vast amount of science done here at the ARL MSRC, we will feature in each issue an interview with some of our more interesting users and scientists. They will share their insights and views of what they do and how they use our HPC and support personnel resources.*

HPC CERT Provides Security for All of DREN

By CERT staff

Since October of 2000, ARL has provided the network intrusion monitoring and detection services for the HPCMP Defense Research and Engineering Network. Known as the HPC CERT, ARL staff are on guard for the Modernization Program 24 hours a day, 7 days a week, observing network traffic and pinpointing malicious behavior. In the two years since ARL has undertaken this mission, there has been a marked improvement in DREN network security and HPC CERT staff have developed an excellent rapport with site security personnel.

National Guard personnel assist in the monitoring duties

One type of report that the HPC CERT generates is called a Poor Security Practice Alert. Poor Security Practices relate to a multitude of actions that could give a prospective hacker an opportunity to gain unauthorized access. Many poor security practices are the result of oversights on the part of legitimate users or system administrators.

Examples of poor security practices include sending a root password in the clear or failing to update software that is known to be vulnerable. For .mil systems, regulations and policies govern the use of the equipment and generally prohibit the use of software not provided and maintained by authorized and trained system administrators.

Sometimes seemingly innocent use of a computer system can result in unexpected consequences. Consider the example in the following paragraphs.

Clockwise from top: Chris Turek, Carlo Mateo and Allan Adkins monitor DREN activity to prevent unauthorized access.

Recently a new employee was issued a computer that had been configured, scanned, and accredited by the local system administrator. Part of the standard software load included a Web browser, as would be expected.

Shortly after settling in with his new computer, the employee inquired about using his system to listen to an Internet "radio" station. He was informed that this was acceptable use of his system, so he searched for a station meeting his listening tastes. The station he selected was a commonly available station and when he chose it, it asked him to accept a browser plugin in order to deliver the service. There was no indication of the function of the plugin during the installation process, and the employee received the radio station as expected when the installation completed.

Within a few hours, the site intrusion detection system was flagging a large amount of network traffic inbound to the employee's computer from many unexpected external sources. Site administration personnel were perplexed by these alerts, but knew it was a new address that they had only recently enabled. When no immediate explanation was found, the system was pulled off the network for detailed examination. It was discovered that the plugin installed by the employee was the culprit.

The plugin was by a company called ChainCast, and it was designed to turn every listener's computer into a rebroadcasting server for additional listeners. Since rebroadcasting was not a service that had been authorized for this computer, this was viewed as a poor security practice by the organization.

Beyond the administrative issues, there are also potential legal ramifications resulting from the use of such software. Two obvious categories of risk are copyright infringement and perceived attacking behavior when acting as a server.

So, the tip for this quarter is to be absolutely certain of the function and purpose of all software installed on your computer. When in doubt, don't install it and contact your system administrator for advice.

New Scientific Visualization Facility at APG

By Rick Angelini and Mark Bolstad

In the Fall of 2002, the ARL MSRC will unveil its latest visualization asset – the Aberdeen Reality Center (ARC), a vision of Dr. N. Radhakrishnan, CISD Director.

The ARC will consist of a 3800-square-foot facility to support both Classified and Unclassified visualization requirements. Existing visualization laboratories will be relocated to this new facility, consolidating resources, and creating an innovative environment for visualization and computational scientists to display their work.

Conceptual rendering of new Reality Center

The centerpiece of the ARC will be a Fakespace RAVE II display system. This system, approximately 40 feet wide by 8 feet tall, features a fixed 20-foot center fixed display panel, and two 10-foot wide reconfigurable "wings" which can be rotated to change the display system from a single, wide wall to a 3-sided cave-like environment. Utilizing the latest generation of Digital Light Processing (DLP) technology, the display will be driven by six Christie Digital Mirage projectors. Each of the projectors in the system will display at a stereo resolution of 1280x1024 pixels, and emit approximately 2000 ANSI lumens of brightness (ten times brighter than traditional CRT-based projection systems), allowing scientists to work in a natural light environment rather than a darkened room. The system is being installed and integrated through the combined efforts of SGI and Fakespace Systems.

The RAVE system will be driven by the latest visualization hardware from SGI, an Onyx 3400. This system will have twenty-four R14K, 500-MHz processors with 24 GB of RAM, and three InfiniteReality3 graphics pipes.

This combination of high-end computational resources and advanced display technology will allow scientists and researchers to experience their data in a highly immersive environment, while enabling the deployment of new visualization techniques such as interactive volume rendering.

The RAVE and its supporting infrastructure have been designed with expandability in mind. As more processing power is required, the SGI system can be easily expanded to accommodate more processors, memory, and graphics pipes.

Additionally, there is room to add different architectures, such as a commodity PC cluster, if the technology proves viable.

When used in conjunction with the existing collaboratorium at the Adelphi Laboratory Center of ARL, a unique collaborative environment will be established. Due to the nature of each facility, large numbers of researchers will be able to interact on complex computational models between the two remote locations.

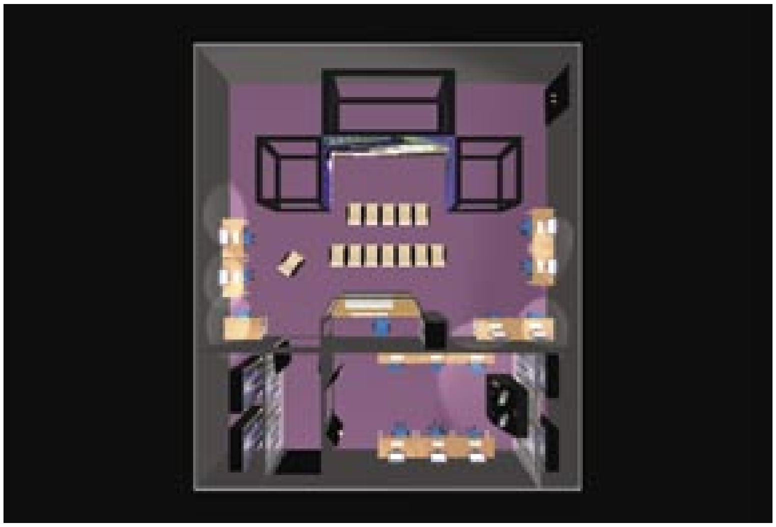

Top view rendering of new Reality Center

A New Approach for Determining Battlefield Parameterization

Dr. Barry A. Bodt, Eric G. Heilman, and Janet F. O'May

The Battlespace Decision Support Team of ARL has explored the applicability of combat simulation to ground command and control (C^2). Specifically, our endeavors have centered on Course of Action (COA) evaluation and metrics for the planning/re-planning process of modern combat. Completely automated COA generation is a difficult proposition because of the complexities and uniqueness of situational combat instances and the elusiveness of capturing subject matter expertise within a programming environment. The lack of an automated COA system led us to examine the advantages of combat simulation to the C^2 process. Consequently, we have developed a statistical methodology to assist commanders in maintaining combat initiative by enhancing future force planning responsiveness and versatility.

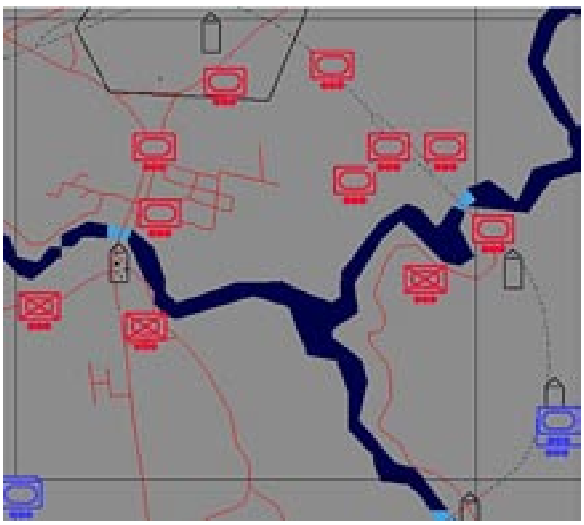

Screen capture of entity placement in OneSAF

The mainstay of our methodology is battlefield parameters. Each parameter is a feature of a battle that is readily available to the participants. We have conducted several experiments designed to improve our understanding of combat simulation capabilities. The basis of our current experimentation is the One Semi-Automated Forces (OneSAF) simulated combat environment. By modifying OneSAF, we created a mechanism that supports the collection of a large combat entity interaction database. Using a two-sided scenario designed to support an even chance for victory to either side, the team gathered a large database from the outcomes of 228 separate scenario executions. While the scenario was small in scale, the sheer number of runs proved to be computationally intensive. The interactive queuing system of the ARL Major Shared Resource Center allowed us to run simultaneous simulations, reducing the amount of time needed to execute the large number of scenarios.

We used several statistical techniques to tap into and identify the meaning of the collected battlefield data. We pared down the large amount of data and focused on 435 parameters. These ranged from the more traditional, such as vehicle damage or attrition rates, to the exotic, such as damage done to a target type by certain munitions during a delimited time period. The 435 parameters included three different time periods during the battle delimited by ammunition expenditure. The first time corresponded to 10% ammunition expenditure, the second time 25%, and the third time 40%. We post-processed the data for use with data mining tools and mined using general linear models, discriminant analysis, logistic regression, and classification and regression trees. The battles averaged 60 minutes in duration. An important result showed that after the first five minutes, the battle outcomes could be classified with 70% accuracy. After 20 minutes, that classification rate rose to 85%.

There are two main conclusions from our experimentation. First, the statistical methodologies used are combat simulation independent. While we chose OneSAF, a training simulation, for our initial studies, the application of statistical methods is not dependent on any particularity of OneSAF for success. Second, as warfare becomes more asymmetric and technology advances become a part of the future force, we must adopt an increasingly flexible planning process to gain and retain combat initiative. Provided we can capture data representing a set of pertinent battle parameters, commanders can use our methodology to generate classifications reflecting battle plan outcomes. Further, the planning staff will gain insights on the importance of tracking non-traditional, time-critical metrics.

The nature of combat will pose unique situations for outcome classification. For instance, the battle might hinge on the presence of a particular unit on a decisive terrain position. In response, we created our methodology with the flexibility to present those metrics that are meaningful in any particular moment during a battle. While the placement of one particular unit may be beyond our scope, the setting of conditions for success is not. For instance, suggesting that more force is necessary on a route inclusive of the key terrain is possible.

Future force planning will rely on new command, control, communications, computers, intelligence, surveillance, and reconnaissance concepts supported by readily available data collection assets. The ability to use high performance computers will enable a reach-back capability. A commander in the field will be able to send a plan to a computer thousands of miles away for analysis. Our methodology will harness the data rich future battlefield to improve the planning/re-planning cycle and assist in maintaining battle initiative.

ARL MSRC Attends ASTC and SMART Conferences

By Brian Simmonds

ARL MSRC personnel attended and participated in both the Advanced Simulation Technology Conference (ASTC) April 14–18 in San Diego, and the Simulation and Modeling for Acquisition, Requirements, and Training (SMART) Conference April 16–18 in Salt Lake City.

Charlie Nietubicz, ARL MSRC Director, acted as session chair of the High Performance Computing Session at ASTC. He organized a panel on "HPC Past, Present, and Future" with center directors from NASA, NSF, and the Army. He spoke on the history of high performance computing centers and their role in solving complex problems.

The SMART Conference was a chance for the MSRC as well as other ARL groups to show some of its more visual efforts. The display booth featured a 3D Immersadesk display with several simulations for the visitors to see. It was also the first outside ARL venue for the Elumens VisionDome, a 3D immersive environment that looks somewhat like a satellite dish. Look for both of these items at future shows. The conference was attended by DoD acquisition personnel as well as members of the scientific community.

User Feedback Aids Customer Service Project

By Dr. Robert Cromp

One of the functions of the ARL MSRC's Applications Support team is to provide answers in a timely manner to users who need assistance in using software packages on the high performance computing facilities. Answering these questions takes priority over all other work. As the size of the user community increases, it is important to develop improvements to this overall process.

Raytheon uses a statistically based methodology known as Six Sigma to develop and improve products and services. Many other companies rely on this same model, introduced originally by Motorola in the late 1980s, including GE, Sony, and Allied Signal. We have previously used the Raytheon Six Sigma process to solicit feedback from the ARL MSRC user community in order to improve upon areas identified explicitly by users. As a follow-on, a recent Raytheon Six Sigma project was conducted to analyze and strengthen our existing approach to handling user support requests.

Once we decided to focus on user support, we needed to prioritize which aspects to address. A method for gathering and grouping ideas known as "affinity diagrams" helped identify two categories: proactively educating the users about the MSRC environment, and organizing and making better use of information.

Of the several ideas generated for proactively educating users about the MSRC environment, two seemed most promising: providing a graphical user interface for automatically generating GRD scripts for the most utilized supported applications, and maintaining frequently asked questions (FAQs) lists for users to consult. The majority of problems users confront involve errors embedded in user scripts for interfacing with the GRD queuing system. These mistakes can be coding problems, improper use of GRD, or incorrect staging of data on the HPC machines. As a result, we have implemented an interface that, through a set of questions, produces a GRD script specifically suited to the user's needs for the applications in Table 1. These applications were chosen due to their high frequency of use at the ARL MSRC. This interface is undergoing pioneer testing and should be available for general use by mid-November.

Table 1: Applications Slated for Automated GRD Script Generation

Abaqus	GAMESS	Overflow
ANSYS	GASP	Patran
Cobalt	Gaussian98	Wind
CTH	ICEMCFD	XPATCH
Fluent	LS-DYNA	

Tom Brezee conducts a customer feedback session at User's Group Conference.

The approach for developing and maintaining meaningful FAQs addresses both categories. To proactively educate users, we are compiling information from the numerous vendors for the various applications on common questions and solutions. This information will be placed on both the Web and within the appropriate application directories. To organize and make better use of information, we performed a comprehensive analysis of the nearly 300 problem tickets answered by the Applications Support team over the past two years. Each ticket was tagged with a set of descriptive keywords, and this information was used to augment the existing Help Desk Oracle database. When a user requests a problem ticket, the ARL MSRC User Support team can now easily retrieve related problem tickets which were previously solved and use these as a basis for formulating a solution to the current problem ticket. In the final phase, we will group the overall set of problem tickets by frequency of keyword occurrence, producing areas where users tend to encounter difficulties. We will then focus on improving user satisfaction through some combination of tool development, user orientation, and improved documentation.

A number of other Raytheon Six Sigma projects are currently in progress, concentrating on areas such as improved large-scale data management methods, techniques for generating complex conceptual animation productions, and educating users on how the ARL MSRC Scientific Visualization team can assist with the visualization of their data sets. All these projects share the common aim of providing high quality customer satisfaction.

Upcoming Events

16 to 22 November 2002 Supercomputing 2002
Location: Baltimore Convention Center, Baltimore, Maryland. Hosted by ARL MSRC.
For more information: http://www.sc2002.org.

2 to 5 December 2002 23rd Army Science Conference
Location: Renaissance Orlando Resort, Orlando, Florida.
For more information: http://www.asc2002.org.

6 to 9 January 2003 41st Aerospace Sciences Meeting and Exhibit sponsored by AIAA
Location: Reno Hilton Resort & Casino, Reno, Nevada.
For more information: http://www.aiaa.org/calendar

PET Events at ARL

CEM Workshop: 10 December 2002
Topic: Use of Time Domain Techniques to Solving High-Frequency Circuit Problems

CCM Training: December 2002
Topic: Advanced Solid-State Modeling

CCM Training: December 2002
Topic: Intermediate/Advanced CASTEP

CDLT Seminar: January 2003
Topic: Computer Distance Learning Technologies Day

CSM Training: January 2003
Topic: Grid Generation and Grid Technology

CEN/CEA Training: January 2003
Topic: Efficient Computational Methods for Circuits and Antennas

The MSRC Help Desk
can be contacted by telephone:
1.800.275.1552
or
410.278.1700

or through e-mail:
msrchelp@arl.army.mil

Contributors

Brian Simmonds, Editor-in-Chief, Photographer, Imaging
Michelle Morgan-Brown, Editor, Writer
Jim Nelson, Photographer
Michael McCraney, Hardware Advisor, Technical Writer

Rick Angelini, ARL
Dr. Barry A. Bodt, ARL
Mark Bolstad, Raytheon
CERT Staff, ARL
Jerry A. Clarke, ARL
Dr. Robert Cromp, Raytheon
Robin Gravels, ARL
Eric G. Heilman, ARL
Robin M. Liles-Meyers, ARL
Pat McKenna, Raytheon
Charles Nietubicz, ARL
Janet F. O'May, ARL
Dr. Steve Wilkerson, ARL

Contacts

Dr. N. Radhakrishnan, Director, Computational and Information Sciences Directorate
radha@arl.army.mil, 410-278-6639

Charles Nietubicz, Director, ARL MSRC
cjn@arl.army.mil, 410-278-3691

Dr. R. "Prabu" Prabhakaran, Raytheon PM
rprabu@arl.army.mil, 410-278-7436

Dr. Andrew Mark, PET Government Lead
amark@arl.army.mil, 410-278-9761

Dr. George K. Lea, Pet II Program Director
glea@hpti.com, 410-278-7195

Thomas Kendall, Customer Service/Chief Engineer
tkendall@arl.army.mil, 410-278-9195

Rick Angelini, Scientific Visualization Team Lead
angel@arl.army.mil, 410-278-6266

Judy Bouchelle-Keithley, Outreach Team Lead
jkeith@arl.army.mil, 410-278-9195

Brian Simmonds, Link Editor-in-Chief
bsimmonds@arl.army.mil, 410-278-3214

Randy Schauer, Webmaster
rschauer@arl.army.mil, 410-278-7533

Helping Hands: Maintenance, Networking, and Facilities Team

By Michelle Morgan-Brown

From left to right, Gary Ruff, Billy Barkaloo, Bob Horner, Mark Motsko, Ken Radke, Israel Magrogan, and Wayne Wood compose the MSRC's hard-working maintenance, networking, and facilities team.

Without our maintenance, networking, and facilities team, the MSRC could not maintain its systems, which means we could not support our user community. In short, the maintenance, networking, and facilities team helps keep this place operational.

Naturally, without a good support team, our mission would not succeed. Happily, the ARL MSRC maintenance, networking, and facilities team, led by Gary Ruff, works hard to keep the center, and therefore our users and their work, up and running every day.

The team is responsible for a variety of functions, including the maintenance of personal computers, workstations, peripheral equipment (printers, scanners, etc.), and low-end SGI and Sun servers.

Ruff's team also oversees the maintenance and tuning of networking routers, switches, and interfaces, while maintaining the site's UPSs and cooling systems for the computing room.

Ken Radke repairs HVAC unit that feeds the computer rooms.

In addition, the team supplies detailed computer room drawings, which are used when assisting in the integration of new systems.

So, next time your home computer crashes, remember that our guys are working hard to make sure that doesn't happen here.

Outreach Team

Welcome to our new format for the Link. It has been a challenge to produce, but we made it. We have had a lot of help and contributions from the staff and users here at ARL MSRC as well as others. We hope that you find the new Link user friendly and helpful. In the future, we will try to bring you not only pertinent information you need, but interesting projects and happenings as well. We are also working on an electronic version of the Link, aptly named eLink. This will be available via our Web site and e-mail subscription.

The team is also feverishly preparing for the upcoming shows. We are the host center for the HPCMO booth at Supercomputing 2002 (SC2002), which is being held in Baltimore this year. We, along with the other MSRCs and DCs, have quite a collection of displays, hardware, and presentations being exhibited this year. Another conference we are going to is the Army Science Conference in Orlando, FL, in early December. There, we also plan to show a host of scientific demonstrations and technologies.

We are pleased to announce the addition of a new team member, Michelle Morgan-Brown. Michelle is our new writer and editor, and we've already put her to work on the Link. Please join us in giving Michelle a warm welcome to the MSRC and the Outreach team.

We would like to thank all of the supporting staff members here at ARL for their help in making the new Link possible with contracting, acquisition, and logistical support. We also would like to give a special thanks to Pat McKenna, who joined us temporarily from Raytheon in Garland, TX, to fill in for our staff writer.

We hope that you like the new format, and we welcome your comments and suggestions. Please feel free to e-mail us at outreach@arl.army.mil.

Enjoy the magazine!

Outreach Team, (l to r) John Vines, Judy Bouchelle-Keithley, Jim Nelson, and Brian Simmonds. Not pictured: Michelle Morgan-Brown.

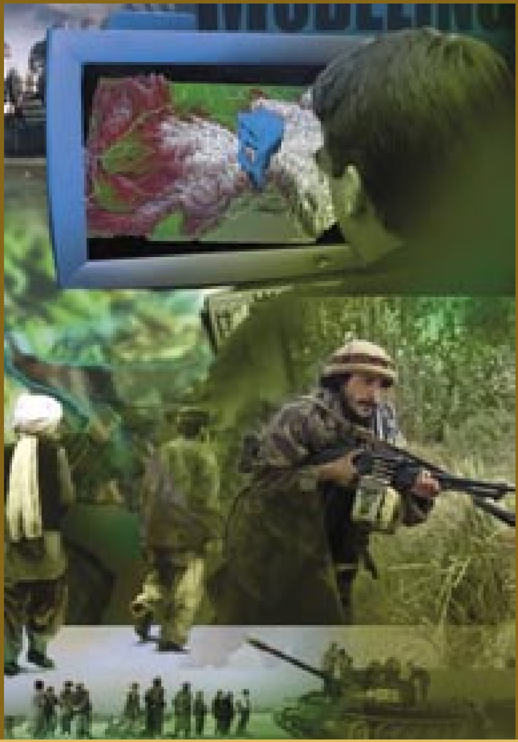

In this Issue:

ARL MSRC Weather Model

Computational Electromagnetics Project: Shilka ZSU23

Austin User's Group Conference Photos

User Interview: Steve Schraml

The MSRC Help Desk can be contacted by
telephone: 1.800.275.1552 or 410.278.1700
or through e-mail:
msrchelp@arl.army.mil

The ARL MSRC effectively supports and empowers DoD researchers with high performance computational techniques, methodologies, and infrastructure

LiNK

Connecting Users to the ARL MSRC

The Army Research Laboratory
Major Shared Resource Center (AMSRL-CI-H)
Aberdeen Proving Ground, MD 21005-5067

STANDARD
U.S. POSTAGE
PAID
APG, MD
PERMIT NO. 1

Visit the ARL MSRC on the web: http://www.arl.hpc.mil